



Self-organization of spatiotemporal receptive fields and laterally connected direction and orientation maps

James A. Bednar*, Risto Miikkulainen

Department of Computer Sciences, The University of Texas at Austin, Austin, TX 78712, USA

Abstract

Studies of orientation maps in primary visual cortex (V1) suggest that lateral connections mediate competition and cooperation between orientation-selective units, but their role in motion perception has not been established. Using a self-organizing model of V1 with moving oriented patterns, we show that (1) afferent weights of each neuron organize into Gabor-like spatiotemporal receptive fields with ON and OFF lobes, (2) these receptive fields form realistic joint direction and orientation maps, and (3) lateral connections develop between patches with similar orientation and direction preferences. These results suggest that a single self-organizing system may underlie the development of orientation selectivity, direction selectivity, and lateral connectivity.

© 2003 Published by Elsevier Science B.V.

Keywords: Direction maps; Orientation maps; Spatiotemporal receptive fields; Direction selectivity; Lagged cells

1. Introduction

Most simple cells in the primary visual cortex (V1) are selective for the direction and orientation of a moving stimulus. Recent measurement techniques have made it possible to plot the neurons' full spatiotemporal receptive fields, which include specific excitatory (ON) and inhibitory (OFF) subregions that vary over time [5]. The functional properties of these cells form a mosaic across V1, with patches of nearby neurons preferring similar directions and orientations [13].

* Corresponding author.

E-mail addresses: jbednar@cs.utexas.edu (J.A. Bednar), risto@cs.utexas.edu (R. Miikkulainen).

In addition to their afferent input from the LGN, the neurons in these maps are connected intracortically through specific long-range lateral connections [7]. The lateral connections have been found to link cells with similar orientation preferences [3,7], which can allow the connections to suppress redundancy in the input and improve the cells' ability to detect changes in a stimulus [2,8,12]. However, the role of these connections in the development and adult function of directional selectivity is not yet clear.

Several computational models have shown that directional selectivity and interleaved orientation and direction maps can develop through activity-dependent self-organization [6,10,14]. However, these simulations have only modeled afferent connection learning, and have treated lateral connections as fixed. In prior work with the LISSOM self-organizing model (Laterally Interconnected Synergetically Self-Organizing Map), we have shown how a Hebbian learning process can develop topographic maps, ocular dominance, orientation, and size preference columns, and patterned lateral connections between them [8,11,12]. We have also shown that these self-organized maps and lateral connections can function in adult visual perception to segment and bind coherent objects and reduce redundancy in the input, and that visual illusions and aftereffects arise through this process [2,4,12]. In this paper we extend the model to develop motion and direction preferences through the self-organization of spatiotemporal receptive fields. Together, these results show that activity-dependent self-organization can explain many of the anatomical and functional characteristics of the cortex.

2. HLISSOM model

These simulations are based on the HLISSOM model [1], which extends LISSOM to include the ON and OFF channels of the LGN. The main idea is that directional selectivity can develop from LGN cells whose output arrives at the cortex after a delay. In the LGN, many cells fire soon after a retinal stimulus. However, other *lagged* cells have recently been found in cat LGN that respond only after a fixed delay [9]. The delay time in these lagged cells varies over a continuous range up to hundreds of milliseconds [15]. We will show below that V1 neurons can use these timing differences to develop spatiotemporal receptive fields.

The architecture for the HLISSOM model is shown in Fig. 1, and will be briefly reviewed below. (For more details, see [2].) The model consists of a hierarchy of two-dimensional sheets of neural units modeling different areas of the visual system: a sheet of retinal photoreceptors, several paired sheets of ON-center and OFF-center LGN units (with a different lag for each pair), and a sheet of cortical units (“neurons”) representing V1. Because the focus is on the two-dimensional organization of V1, each neuron corresponds to a vertical column of cells through the six anatomical layers of the cortex.

The model simulates behaviour at three different timescales, denoted by variables f , t , and s in the equations below. Variable f represents a single fixation out of a sequence of thousands presented to the model. At each fixation, motion in the image

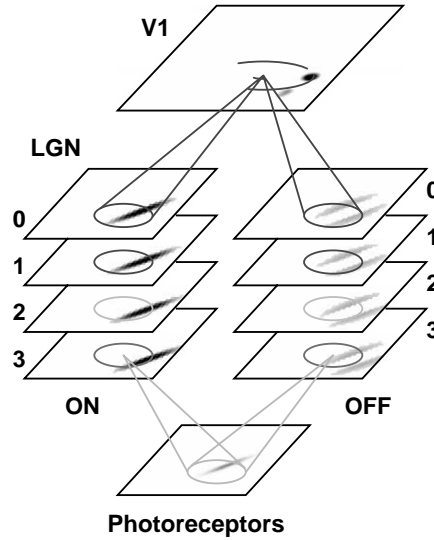


Fig. 1. HLISSOM model of orientation and direction selectivity. The model is a hierarchy of sheets of neural units, modeling the early visual pathway. In this implementation, there are eight sheets of LGN units. The LGN units receive input from their circular receptive fields (RFs) on the photoreceptors. Connections to a sample photoreceptor RF are plotted for two LGN units; all units at the same position in the other six LGN sheets also receive input from that same RF. Each LGN sheet contains a homogenous group of units, with either ON- or OFF-center organization and a lag of 0, 1, 2, or 3 timesteps. Similarly, neurons in V1 have circular RFs on all eight LGN sheets; connections to two of the RFs are shown. V1 neurons also have lateral excitatory (small dotted circle) and lateral inhibitory (large dashed circle) connections to their neighbors. Moving input patterns are drawn on the photoreceptor sheet in discrete timesteps, like frames of a movie. At the first timestep, the ON and OFF LGN cells with time lag 3 compute their activity. At each subsequent timestep, the input pattern is moved slightly and LGN cells with lags 2, 1, and 0 compute their activity in turn. Once all LGN cells have been activated, initial V1 activity is computed from the LGN responses, and the activity then spreads laterally within V1. Both afferent and lateral V1 weights are adapted when the activity settles.

is simulated at several discrete timesteps t . For a given time t , lateral interactions in the cortex are computed for several timesteps s .

The input to the model is a series of activity patterns on the sheet of photoreceptors, such as two-dimensional oriented Gaussians or grayscale photographic images. A sample input is drawn in the photoreceptor sheet of Fig. 1. For a given moving input, the photoreceptor activity is drawn at multiple discrete timesteps. At each such time step t , the activity levels of all LGN cells with lag t are calculated. Each LGN cell (i, j) with lag t computes its response η_{ij} as a scalar product of a fixed weight vector and its receptive fields on the photoreceptor sheet at time t :

$$\eta_{ij} = \sigma \left(\sum_{\rho ab} \gamma_{\rho} X_{\rho ab} w_{ij, \rho ab} \right), \quad (1)$$

where ρ identifies the input sheet (the photoreceptors, in this case), σ is a piecewise linear sigmoid activation function, γ_ρ is a constant scaling factor, $X_{\rho ab}$ is the activation of input unit (a, b) on sheet ρ at timestep t , and $w_{ij, \rho ab}$ is the corresponding weight value. Each V1 neuron computes its initial response like that of an LGN cell, except that ρ identifies one of the ON or OFF LGN sheets. After the initial response, the V1 activity settles through short-range excitatory and long-range inhibitory lateral interaction:

$$\eta_{ij}(s) = \sigma \left(\sum_{\rho ab} \gamma_\rho X_{\rho ab}(s-1) w_{ij, \rho ab} \right), \quad (2)$$

where ρ identifies either an LGN sheet or the lateral excitatory or inhibitory weights to V1, γ_ρ is a constant scaling factor for each ρ (negative for inhibitory lateral weights), and $X_{\rho ab}(s-1)$ is the activation of input unit (a, b) during the previous settling step. The V1 activity pattern starts out diffuse, but within a few iterations of Eq. (2), converges into a small number of stable focused patches of activity, or activity bubbles. After the activity has settled, the connection weights of each V1 neuron are modified. All V1 weights adapt according to the Hebb rule, normalized so that the sum of the weights from each sheet ρ is constant for each neuron (i, j) :

$$w_{ij, \rho ab}(f+1) = \frac{w_{ij, \rho ab}(f) + \alpha_\rho \eta_{ij} X_{\rho ab}}{\sum_{ab} [w_{ij, \rho ab}(f) + \alpha_\rho \eta_{ij} X_{\rho ab}]}, \quad (3)$$

where η_{ij} stands for the activity of neuron (i, j) in the final activity bubble, $w_{ij, \rho ab}(f)$ is the connection weight from the previous fixation, α is the learning rate for each type of connection, and $X_{\rho ab}$ is the presynaptic activity. The larger the product of the pre- and post-synaptic activity $\eta_{ij} X_{\rho ab}$, the larger the weight change. At long distances, few neurons have correlated activity and therefore most long-range connections eventually become weak. The weakest connections are eliminated periodically, resulting in patchy lateral connectivity similar to that observed in the visual cortex.

For the experiments reported in this paper, four 36×36 ON-center and four 36×36 OFF-center cell sheets received input from a 54×54 photoreceptor sheet. Each ON/OFF cell had a fixed Difference of Gaussians receptive field (RF) within the photoreceptor array. Initially, the afferent weights of the 142×142 V1 neurons were random, and the lateral weights had a smooth circular Gaussian profile. The learning parameters were the same as in our earlier V1 orientation model [2], scaled for this cortex size using the model scaling methodology presented in [1].

3. Results

Fig. 2 shows the self-organized afferent weights for a representative neuron after 20,000 image fixations. Nearly all neurons developed spatiotemporal receptive fields selective for orientation, and most were also selective for direction. These RFs are similar to those found experimentally in the cortex [5]. As in V1 cells, the model RFs




Fig. 2. Self-organized spatiotemporal receptive field. These images show the afferent weights for a typical neuron, plotted on the retina. As a visualization of the preferred stimulus of this neuron, the OFF weights were subtracted from the ON for each lag. In the resulting plots, white indicates a net ON subregion of the retina, i.e. an area that will excite the neuron if it is brighter than its surround. Black indicates OFF subregions. Together, the afferent weight plots show that the most effective stimulus for this neuron would be a horizontal bar or grating moving upwards. More specifically, this neuron will be highly active at time t if there was a bright bar aligned with the ON subregion in (a) at time $t - 3$, a bright bar aligned with the ON subregion of (b) at time $t - 2$, and so on until the present time t , where the bar will be aligned with the ON subregion of (d). Visual cortex neurons in animals have similar spatiotemporal properties [5].

can only rarely be expressed as a function of a fixed spatial RF and a time-dependent scaling factor. The RFs are thus *space-time inseparable* [5]. Also as in the cortex, the orientation and direction preferences of a neuron were generally perpendicular.

Fig. 3 shows that the self-organized global orientation and direction maps are also similar to those found in animals [13]. For instance, a patch of neurons highly selective for one orientation and direction of motion will often have an adjacent or contiguous patch selective for the same orientation, but the opposite direction. Long-range lateral connections within the map follow this global organization, connecting neurons similar in both orientation and direction preference. These connection patterns represent the correlations between units over the course of self-organization. As will be described below, the self-organized lateral weight patterns are likely to play a crucial role in adult visual perception as well as development.

4. Discussion and future work

The results demonstrate that a single Hebbian learning algorithm can explain how topographic maps, spatiotemporal receptive fields, and lateral connections synergistically self-organize from moving stimuli. The model predicts that long-range lateral connections in V1 will be found to connect neurons with similar direction preference as well as orientation preference. In future work we will validate the model against experimental results from particular species and study how the orientation, direction preference, and lateral connectivity interact with other stimulus dimensions, such as spatial frequency preference.

Other studies will investigate how the direction map and its connections operate in the adult. The self-organized connection patterns should allow moving objects to be bound together and segmented from other objects, as previously shown for static




Fig. 3. Self-organized orientation and direction map. Through self-organization, a smoothly varying map of orientation and direction preferences emerged. Each of the 142×142 neurons is colored with its preferred orientation, according to the key at right. The orientation map is similar to those found experimentally in animals, containing features such as pinwheels, linear zones, and fractures [13]. The direction preferences for every third neuron are plotted as white arrows overlaid on the orientation map. The size of each arrow is proportional to the neuron's directional selectivity, and the arrow points in the neuron's preferred motion direction. The preferred orientation is generally perpendicular to the preferred direction. As an example, the neuron from Fig. 2 is marked with a black square in the lower center of the plot. The neurons in that region of the map prefer horizontal stimuli moving upwards. Neurons in the patch just below the marked neuron also prefer nearly horizontal lines, but moving in the opposite direction. Other paired patches can be seen elsewhere in the map, as found experimentally [13]. The black outline shows the extent of the self-organized lateral weights for the marked neuron. The short-range connections target neurons with a variety of preferences, but long-range connections (a) target similar orientations and direction preferences, (b) extend along the orientation axis (horizontal), and (c) avoid orthogonal orientations and opposite directions. Animations of this plot are available at www.cs.utexas.edu/~nn.

patterns [4]. Short-term direction-specific adaptation of the lateral connections should replicate the motion aftereffect, also called the waterfall illusion, as we found for the tilt aftereffect [2]. This work will help unify explanations of development and adult visual function into a coherent theory of the visual cortex.

5. Conclusion

Hebbian learning of moving, oriented patterns in a self-organizing model can explain how orientation selectivity, direction selectivity, and lateral connectivity develop synergetically in V1. The model provides specific predictions for the role of the lateral connections in the development and function of the direction map. This study and future work will help strengthen our understanding of the visual cortex as a continuously adaptive, self-organizing system.

Acknowledgements

Supported in part by the National Science Foundation under grants IRI-9309273 and IIS-9811478 and by the National Institutes of Mental Health under Human Brain Project grant 1R01-MH66991.

References

- [1] J.A. Bednar, Learning to See: Genetic and Environmental Influences on Visual Development, Ph.D. Thesis, Department of Computer Sciences, The University of Texas at Austin, 2002.
- [2] J.A. Bednar, R. Miikkulainen, Tilt aftereffects in a self-organizing model of the primary visual cortex, *Neural Comput.* 12 (7) (2000) 1721–1740.
- [3] W.H. Bosking, Y. Zhang, B. Schofield, D. Fitzpatrick, Orientation selectivity and the arrangement of horizontal connections in tree shrew striate cortex, *J. Neurosci.* 17 (6) (1997) 2112–2127.
- [4] Y. Choe, R. Miikkulainen, Self-organization and segmentation in a laterally connected orientation map of spiking neurons, *Neurocomputing* 21 (1998) 139–157.
- [5] G.C. DeAngelis, I. Ohzawa, R.D. Freeman, Receptive-field dynamics in the central visual pathways, *Trends Neurosci.* 18 (1995) 451–458.
- [6] I. Farkas, R. Miikkulainen, Modeling the self-organization of directional selectivity in the primary visual cortex, *Proceedings of ICANN-99*, Springer, Berlin, 1999, pp. 251–256.
- [7] C.D. Gilbert, T.N. Wiesel, Columnar specificity of intrinsic horizontal and corticocortical connections in cat visual cortex, *J. Neurosci.* 9 (1989) 2432–2442.
- [8] R. Miikkulainen, J.A. Bednar, Y. Choe, J. Sirosh, Self-organization, plasticity, and low-level visual phenomena in a laterally connected map model of the primary visual cortex, in: Goldstone, et al. (Eds.), *Perceptual Learning*, Vol. 36 of *Psychology of Learning and Motivation*, Academic Press, New York, 1997, pp. 257–308.
- [9] A.B. Saul, A.L. Humphrey, Evidence of input from lagged cells in the lateral geniculate nucleus to simple cells in cortical area 17 of the cat, *J. Neurophys.* 68 (4) (1992) 1190–1208.
- [10] H. Shouno, K. Kurata, Formation of a direction map by projection learning using Kohonen's self-organization map, *Biol. Cybernetics* 85 (4) (2001) 241–246.
- [11] J. Sirosh, R. Miikkulainen, Cooperative self-organization of afferent and lateral connections in cortical maps, *Biol. Cybernetics* 71 (1994) 66–78.
- [12] J. Sirosh, R. Miikkulainen, J.A. Bednar, Self-organization of orientation maps, lateral connections, and dynamic receptive fields in the primary visual cortex, in: Sirosh, et al. (Eds.), *Lateral Interactions in the Cortex: Structure and Function*, Austin, TX, The UTCS Neural Networks Research Group, ISBN 0-9647060-0-8, www.cs.utexas.edu/~nn/web-pubs/htmlbook96, 1996.
- [13] M. Weliky, W.H. Bosking, D. Fitzpatrick, A systematic map of direction preference in primary visual cortex, *Nature* 379 (1996) 725–728.

- [14] S. Wimbauer, O.G. Wensch, J.L. van Hemmen, K.D. Miller, Development of spatiotemporal receptive fields of simple cells: II, *Biol. Cybernetics* 77 (6) (1997) 463–477.
- [15] J. Wolfe, L.A. Palmer, Temporal diversity in the lateral geniculate nucleus of cat, *Visual Neurosci.* 15 (4) (1998) 653–675.



Jim Bednar is a postdoctoral researcher at the University of Texas at Austin. He completed his Ph.D. at UT in spring 2002, and also has an M.A. in Computer Science, a B.A. in Philosophy, and a B.S. in Electrical Engineering. His research focuses on computational modeling of cortical map development. Jim is the lead author of the forthcoming Topographica modeling software package, under development through a Human Brain Project grant from the National Institutes of Health.



Risto Miikkulainen is a Professor of Computer Sciences at the University of Texas at Austin, and is on the faculty of the Institute for Neuroscience and the Center for Perceptual Systems. He received his Ph.D. in Computer Science from UCLA in 1990, and his M.S. in Applied Mathematics from the Helsinki University of Technology in 1986. His research focuses on applications of neural networks and genetic algorithms to cognitive science, computational neuroscience, and neuroevolution.